On the structure of learning agents

Alok Raj

July 13, 2012

Abstract

This paper presents the thesis that all learning agents of finite information size are limited by their informational structure in what goals they can efficiently learn to achieve in a very complex environment. The thesis implies that there is no efficient universal learning algorithm. An agent can go past the learning limits imposed by its structure only by slow evolutionary change or blind search which in a very complex environment can only give an agent an inefficient universal learning capability that can work only in evolutionary timescales or improbable luck.

1 Introduction

Can a computing machine learn to achieve all goals in a complex environment? How complex a machine needs to be in order to do so? In this paper we attempt to define and ask these questions and discuss a possible answer, and the constraints and limitations the answer implies for all learning machines or algorithms.

Machine learning is a widely studied area. In most cases machine learning deals with supervised or assisted learning where information about the environment is passed to a learning machine for training during the learning phase, for example, a list of labeled examples of patterns for learning a concept. Assisted or supervised learning is a mature field and good progress has been made over the years, for example, back propagation training for multilayer neural networks [RHW86], support vector machines [VC95], and Probabilistically Approximately Correct (PAC) model of computational learning proposed by L. Valiant [Val84].

Discovering which patterns are important for achieving a certain goal in an unknown environment is an important learning problem in its own right. Machine learning also studies unsupervised learning, for example,

for neural networks to learn the statistical structure in input data without training examples [HS99], or for clustering of unlabeled data, or for feature extraction or dimensionality reduction for large input datasets, which are steps in this direction.

In this paper, we study unassisted learning in an unknown environment, where a learning machine does not receive any information about its environment from an external agent. Henceforth in this paper, by learning we always mean unassisted learning in this sense.

2 A Learning Agent Abstraction

To achieve goals in an environment, a machine needs a body with sensors to sense and actuators to manipulate its environment. We use the term "agent-body" for this, and assume that the agent-body has sufficient sensing and physical action capabilities for achieving a sufficiently large set of goals in a complex environment to allow a meaningful discussion of goal achieving and learning capabilities. Defining exact specifications of the agent-body is not relevant to the thesis presented here except that when we compare the goal achieving capabilities of different machines, we assume that they all have physically identical agent-bodies.

A machine also requires a computing platform part and an information processing algorithm to be able to achieve goals in a complex environment. We define and use the term "agent" as the algorithm or the Turing machine (both used equivalently here) running on the computing platform part which processes all the incoming sensory information and controls all the physical actions of the agent-body. Any data an agent may acquire during its exploration of the environment is considered a part of the agent. We consider two agents A and B identical only when a binary string[1] representation for A including all its data is identical to that for agent B under the same representation scheme. The *size* of an agent A, denoted by $size(A)$, is defined as the size in number of bits for the binary string representation for A including all its data. We assume an agent to have as much blank memory available for its use as it requires. An extension of the Church-Turing Thesis [Tur48] is premised that states that the learning and goal achieving capabilities of any physically realizable system can be implemented by a Turing machine running on a computing platform of comparable throughput, controlling an agent-body which is physically equal in terms of physical action capabilities and sensory capabilities. This is equivalently stated by the physical symbol

[1] a bit string, i.e., a member of $\{0,1\}^*$.

system hypothesis of Newell and Simon [NS76] (also see [Nil07] for a recent analysis of its current status):

"A physical symbol system has the necessary and sufficient means for general intelligent action."

We define a *goal* to be a pair of environment states (s_1, s_2), where s_1 is the starting state and s_2 is the end state. The physical state of the agent-body outside the computing platform part is considered to be a part of the environment state.

A *goalset* is a set of goals $\{(s_i, s_j) \mid s_i, s_j \ \textit{are environment states}\}$. We define ALL_GOALS for an environment S to be the set of all goals on S that the agent-body can physically achieve.

An agent achieves a goal $g = (s_1, s_2)$ using some sequence of actions that takes the environment through a path, that is, a sequence of states, which has cost, for example, energy cost, computational cost, environmental cost, and possibly other costs. We represent the sum of all the costs for achieving a goal g through some path by a positive real number and denote it by $cost(g)$. We note that $cost(g)$ can be different for different paths used for the same goal g. Different paths for the same goal g can result from the agent not repeating the same sequence of actions, or from the randomness in the environment not resulting in the same state transitions for the same actions.

To determine if an agent knows how to achieve a goal $g = (s_1, s_2)$, we arbitrarily set a "target cost" for each goal using the function[2] $pcost : ALL\_GOALS \rightarrow R^+$, and say that an agent A knows how to achieve a goal g, or use the phrase "an agent A can *effectively* achieve a goal g" if and only if the agent A can achieve the goal g within $median(cost(g)) \leq pcost(g)$ in repeated trials with identical start state s_1 and identical agent A at start.

When we say "an agent A can effectively achieve a goalset G" we mean $\forall g \in G$, *A can effectively achieve g.*

By exploring its environment, an intelligent agent can learn to effectively achieve a goal g even if it can not effectively achieve the goal g to start with. Let us denote the cost for all the explorations by an agent A for learning to effectively achieve a goal $g = (s_1, s_2)$ starting from state s_1 by $learning\_cost(g, A)$. This cost would depend on g, as well as on the agent

[2]choice of pcost() is arbitrary and immaterial for the thesis presented here, and we can choose the value of $pcost(g)$ appropriately approaching the optimal cost for each g to indicate the agent's goal achieving capability within a certain cost with any good enough success rate.

A and the state of its knowledge about its environment. Also, this cost may vary if we repeat the experiment to find $learning\_cost(g, A)$ under identical start state s_1 and identical agent A.

To determine if an agent A can *efficiently*[3] learn to effectively achieve a goal $g = (s_1, s_2)$, we arbitrarily set a "target learning cost" for each goal using the function $lcost : ALL\_GOALS \rightarrow R^+$, and use the phrase "agent A can efficiently learn to effectively achieve goal g" or "agent A can efficiently learn goal g" if and only if $median(learning\_cost(g, A)) \leq lcost(g)$ in repeated trials with identical start state s_1 and identical agent A at start.

When we say "an agent A can efficiently learn a goalset G" we mean $\forall g \in G$, *A can efficiently learn g.*

3 Principles of Learning

An agent needs information about its environment to be able to effectively achieve goals in a complex environment. We hypothesize that there is no finite sized body of information that would allow an agent to effectively achieve any arbitrary goalset in an environment of unbounded complexity. That is, every finite sized and fixed agent has a fixed largest goalset that it can effectively achieve. An agent can learn and thereby change itself to extend the largest goalset it can effectively achieve. We state this more formally in the following principle.

Principle of Fixed Goal Domain. For every finite sized agent A in an environment S, there exists a fixed largest goalset
$GOAL\_DOMAIN(A) \subset ALL\_GOALS$ that A can effectively achieve.

We note that many different agents A_i of arbitrarily large sizes can have the same $GOAL\_DOMAIN(A_i)$.

We hypothesize that an agent would need more information about its environment to achieve more goals. As a result, the size of the smallest agent that can effectively achieve a given goalset can be no smaller than a certain critical size, and this size would grow with larger (superset) goalsets requiring more environmental information. The phrase "effectively achieve" as defined above is important here. It is possible for a smaller sized agent which does not have the required environmental information to randomly achieve goals at low cost sometimes, but without repeatable performance.

[3]the term "efficient" is used here differently than the usage of the term to mean polynomial computation time in computational complexity theory.

We define the term $Description\_Complexity(G)$ for a goalset G in an environment S as an integer value, such that for all agents A_i that can effectively achieve G, $size(A_i) \geq Description\_Complexity(G)$[4].

We state the second principle below.

Principle of Increasing Complexity. $Description\_Complexity$ grows monotonically with larger (superset) goalsets. That is, for all goalsets $G_1, G_2 \subseteq ALL\_GOALS$,

$$G_2 \supseteq G_1 \rightarrow Description\_Complexity(G_2) \geq Description\_Complexity(G_1).$$

We note that only for trivial superset goalsets requiring no new environmental information $Description\_Complexity$ would not grow. For superset goalsets requiring new environmental information the $Description\_Complexity$ would be larger. We will use the phrase "*Description* of environment S relative to a goalset G", or in short "*Description*" when S and G are clear from the context, to denote an agent of size $Description\_Complexity(G)$ that can effectively achieve G. We note that the smallest agent for a goalset G may not be unique. In other words, there can be multiple alternative *Descriptions* of an environment S relative to a goalset G.

Additionally we need one more principle for learning that all agents would be constrained by. We hypothesize that an agent cannot efficiently learn and build a large *Description* of its environment starting with zero information about its environment. An agent must have the right environment representational building blocks, or what we will call $Microconcepts$ henceforth, to be able to efficiently learn a required *Description* of its environment to be able to effectively achieve goals. Again, the phrase "efficiently learn" as defined above is important here. It is conceivable to have an algorithm for evolutionary change, with zero information about its environment to start with, that can learn to achieve a large goalset, but only over a very large evolutionary timescale. Here the term evolutionary change means a change that is not a pre-calculated change made for achieving a known benefit. The benefit of change towards the knowledge of the environment and

[4]Note that $Description\_Complexity(G)$ is not the same as Kolmogorov complexity for an agent that can effectively achieve G as its largest goalset. Agents that an effectively achieve the same goalset G as their largest goalset and their Kolmogorov complexity can be arbitrarily large, as an agent can include an arbitrarily long and useless random string.

the target goalset can only be evaluated by experimenting in the environment after the change is made, since the agent does not have the required knowledge of the environment for an evaluation beforehand. We also assume that the luck of any physical system is limited by the laws of probability. In the case of agents, this means that the probability of success of generating a large Description bit string by random guessing would be exponentially diminishing with the increasing size of the solution string. Exhaustive blind search for larger Description strings would have exponentially larger cost. The third principle states this hypothesis below.

Principle of Microconcepts for Learning Agents. For an environment S and a goalset G on S, an agent must have, at the start, a set of environment specific bit strings, henceforth called $Microconcepts(G)$, from a class of possible such sets, to be able to efficiently learn the goalset G. For any agent A and a goalset G' for which A does not have a complete required set for a possible $Microconcepts(G')$, A would be limited to an "inefficient" learning through evolutionary change to learn the missing members of $Microconcepts(G')$ where the cost would grow exponentially with the size of the change required, before it can have the capacity to efficiently learn G'.

Let us denote the size of the smallest agent that contains a $Microconcepts(G)$ that allows it to efficiently learn a goalset G by $Critical\_Agent\_Size(G)$. We extend the principle of increasing complexity to learning agents below.

Principle of Increasing Complexity for Learning Agents. Agents that can efficiently learn larger (superset) goalsets requiring newer *Microconcepts* would be larger in size. That is, for all goalsets $G_1, G_2 \subseteq ALL\_GOALS$, such that $G_2 \supset G_1$, and all possible $Microconcepts(G_2) \backslash Microconcepts(G_1)$ are non empty sets,

$$Critical\_Agent\_Size(G_2) > Critical\_Agent\_Size(G_1).$$

4 Consequences

Above principles allow us to have a notion of complexity of an environment S relative to a goalset G on S, given by $Description\_Complexity(G)$ and $Critical\_Agent\_Size(G)$. Based on these we can define two kinds of infinitely complex environments.

The first kind of infinitely complex environment is one where $Description\_Complexity(ALL\_GOALS)$ is infinitely large, but $Critical\_Agent\_Size(ALL\_GOALS)$ is finite in size.

The second kind of infinitely complex environment is one where both $Description\_Complexity(ALL\_GOALS)$ and $Critical\_Agent\_Size(ALL\_GOALS)$ are infinite in size.

A direct implication of the above principles is that all finite sized agents in an infinitely complex environment would be domain limited in the largest goalset they can effectively achieve. For every finite sized agent, there would be goals that an agent cannot achieve effectively even though the agent-body would be physically capable of effectively achieving them. The largest goalset a finite sized agent can effectively achieve would be a proper subset of ALL_GOALS.

As a consequence of the principles for learning agents, there will be a Goalset $G' \subset ALL\_GOALS$ that any finite sized agent cannot efficiently learn in an infinitely complex environment of the second kind. The learning boundaries of an agent, or what goals an agent can efficiently learn in a very complex environment, is limited by the informational structure the agent has. More capable agents would have a larger environment specific informational structure to allow them to efficiently learn larger goalsets. Slow evolutionary change is one way for an agent to achieve its learning capacity. Another way is that an agent can copy the environment representational information from another, already evolved agent (which implies information transfer from another agent).

Consider an agent of a finite size in an environment S of unbounded complexity. Let us assume hypothetically that this agent has universal general intelligence, that is, has the ability to efficiently learn any arbitrary goalset $G \subset ALL\_GOALS$ on S. Such an agent would directly violate the above principles for learning agents.

The thesis presented here also implies that no brain simulation where a small sized cortical model is replicated in large number, and as a result the model compressible to a small size, and where the model is uninformed by the environment specific information, i.e., the microconcepts, encoded in the human or animal brain given by evolution, would succeed in replicating the human or animal learning and goal achieving capabilities.

5 Conclusion

We have presented the thesis that more capable agents have larger complexity, or more precisely, the smallest agents that can achieve or learn to achieve more goals in a complex environment have larger complexity (in size).

We have proposed that for efficient learning, a machine must start with pre-encoded knowledge of its environment, or *Microconcepts*, which decide the efficient learning boundaries for a machine. A machine can only learn beyond its efficient learning boundaries at a much slower pace using evolutionary methods. This also discounts the possibility where we would cross a certain threshold in AI development which would allow an AI system to recursively create increasingly more powerful AI systems in a very short span of time, thereby creating a system of much larger intelligence.

5.1 Future work

Are the principles presented here derivable as a consequence of $P \neq NP$ (assuming $P \neq NP$) [Aar11, Wig07]?

Are there goals in an environment that are physically achievable by an agent-body but computationally not effectively achievable by any agent? Can uncomputable problems be mapped to physical goals?

6 Acknowledgements

I would like to thank the web for providing easy access to the vast amount of work done in this field and to all the people on whose work this work builds on.

References

[Aar11] S. Aaronson. Why Philosophers Should Care About Computational Complexity. to appear in Computability: Gödel, Turing, Church, and Beyond, edited by B. J. Copeland, C. Posy, and O. Shagrir, MIT Press, 2012.

[HS99] G. E. Hinton and T. J. Sejnowski (editors) (1999): Unsupervised Learning: Foundations of Neural Computation, MIT Press, ISBN 0-262-58168-X, 1999.

[Nil07] N. Nilsson. The Physical Symbol System Hypothesis: Status and Prospects. M. Lungarella, et al., (eds.), 50 Years of AI, Festschrift, LNAI 4850, pp. 9-17, Springer, 2007.

[NS76] A. Newell and H. A. Simon. Computer Science as Empirical Inquiry: Symbols and Search. Communications of the ACM, 19(3):113-126, 1976.

[RHW86] D. E. Rumelhart, G. E. Hinton and R. J. Williams. Learning representations by back-propagating errors. Nature 323 (6088): 533536. DOI:10.1038/323533a0, 1986.

[Tur48] A. M. Turing. Intelligent Machinery. National Physical Laboratory Report (1948), in B. Meltzer and D. Michie, eds. Machine Intelligence 5, Edinburgh University Press, 1969.

[Val84] L. G. Valiant. A theory of the learnable. Communications of the ACM, 27(11):1134-1142, 1984.

[VC95] V. Vapnik and C. Cortes. Support Vector Networks. Machine Learning, vol. 20, pp. 273-297, 1995.

[Wig07] A. Wigderson. P, NP and Mathematics - a computational complexity perspective. Proceedings of the ICM 06 (Madrid), vol. 1, EMS Publishing House, Zurich, pp. 665-712, 2007.